

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

Virtualization Of I/O Adapter Resources

Background of Invention

[0001] FIELD OF THE INVENTION

[0002] The subject invention relates to hardware-to-hardware data transmission in computer systems. In particular, it relates to method and system for operating I/O adapters attaching computing devices either to an I/O periphery, a network, or other computing devices.

[0003] DESCRIPTION AND DISADVANTAGES OF PRIOR ART

[0004] The area of the invention concerns hardware of computer systems and network components. It deals more particularly with a method to improve the performance of I/O adapters and utilisation of adapter-local resources like memory.

[0005] As revealed by a first publication of the InfiniBand Architecture (IBA), the prior art of interconnect technologies have failed to keep pace with the current computer evolution and the increased burden imposed on data servers, application processing and enterprise computing created by the popular success of the Internet.

[0006] High end computing concepts such as clustering, fail-safe operations, and 24x7 hour availability demand greater capacity to move data between processing nodes as well as between a processor node and I/O devices. These trends require higher bandwidths and lower latencies, they are pushing more functionality down to the I/O adapters, and they are demanding greater protection, higher isolation, deterministic behavior, and a higher quality of service than it is currently available. InfiniBand helps to achieve the above mentioned aims.

[0007] The invention can be advantageously applied with this new InfiniBand technology and thus increases speed of technical evolution.

[0008] Although the invention has a quite general scope it will be discussed and set out with reference to a specific prior art hardware-to-hardware data transmission in computer systems. This is a communication between a CPU subsystem 8 and a host adapter 18, as depicted in fig.1 and explained next below.

[0009] Today's computer systems (hosts) have a "dense-packed" CPU-memory-subsystem 8 comprising a plurality of CPUs with caches 10, system memory 12, memory controller 14, interconnect logic, etc. Input/output devices, further referred to herein as I/O devices 16 like storage devices, communication networking devices, inter-system connections, etc. are attached via a so-called I/O or host adapter 18. The host adapter 18 may be connected with some "distance" in terms of access time to the CPU-memory subsystem.

[0010] Applications running in the CPUs use specific communication protocols for their connections to said I/O devices 16 and other computer systems accessible via a network.

[0011] These protocols, as for example InfiniBand mentioned above, may define that the application can post work requests to the system memory and is enabled to signal the host adapter to process these work requests. This requires, however, that for signaling and control purposes some amount of information has to be transferred from the CPU-memory-subsystem 8 to the host adapter 18. There are protocols which define very complex tasks for the host adapter to execute in order to perform said processing of the work requests. As it is apparent to a person skilled in the art, a multiple queue processing system is used for processing various incoming requests, in-/outbound data traffic associated with work queues, and system control queues.

[0012] In prior art there have been two different types of methods to cope with this problem:

[0013] With the first type of methods, the I/O adapter 18 is equipped with local memory 20, e.g. implemented on-chip or as separate SRAM/DRAM on card or board. The required control information of the posted work requests is stored in this local

memory. During processing, the host adapter 18 has fast access to the required information. This approach performs very well, but there are resource restrictions, for example the relatively small maximum number of postable work requests which prevents this prior art approach from scaling up to larger environments. This is primarily due to size limitations of the local memory 20. A simple up-scaling of the local memory is expensive as it costs too much (e.g. chip area costs or SRAM/DRAM module costs).

[0014] With the second type of methods, the I/O adapter is not equipped with local memory. Instead, it contains a small set of registers in logic to hold the required control information of one or more work requests. Processing work requests requires many accesses to system memory. This approach is optimized for cost but would not perform well although it does not imply the resource restrictions of method 1. This approach would be a significant obstacle for implementing a well performing, fabric-based switching technology such as InfiniBand.

[0015] It is thus an objective of the invention to overcome the performance / resource restriction problems as outlined above while concurrently being compatible with the switching technology in general.

Summary of Invention

[0016] These objects of the invention are achieved by the features stated in enclosed independent claims to which reference should now be made. Further advantageous arrangements and embodiments of the invention are set forth in the respective subclaims.

[0017] According to a primary aspect of the invention a method for improving the performance of a network coupling adapter is disclosed which attaches one or more computing devices via an interconnected memory, to either one of an I/O periphery, a network, or other computing devices. The method is then characterized by the steps of: operating a local memory being associated with the network coupling adapter as a cache memory relative to a system memory, called an interconnected memory, associated with one or more computing for storing transmission control information.

[0018] Various other objects, features, and attendant advantages of the present invention

will become more fully appreciated as the same becomes better understood when considered in conjunction with the accompanying drawings, in which like reference characters designate the same or similar parts throughout the several views.

Brief Description of Drawings

- [0019] Fig. 1 is a schematic block diagram showing the structural elements in a prior art computer system being equipped with a host adapter.
- [0020] Fig. 2 is a schematic block diagram showing the basic structure of an inventional method for caching queue pairs in a first operating state.
- [0021] Fig. 3 is a schematic block diagram according to fig. 2 in a second operating state different to that one shown in fig. 2.
- [0022] Figs. 4A and 4B are a schematic diagram showing the basic steps of the associated control flow according to the inventional embodiment shown in figs. 2 and 3.

Detailed Description

- [0023] The embodiment described next is directed to a design for a transport layer logic of an I/O adapter, i.e., a so-called Host Channel Adapter (HCA) as defined by the InfiniBand Architecture.
- [0024] The term network and network coupling adapter is to be understood in a very general sense: The network can be for example a Wide Area Network (WAN), a Local Area Network (LAN), or even backplane bus within a PC where the bus participants are interpreted as network attached elements. The term network coupling adapter is thus any hardware device in such hardware structure which interconnects network components.
- [0025] Said hardware structures include explicitly so-called fabric structures as well as a replacement technology of any kind of conventional bus technology. The expression "fabric" has the general meaning of 'configuration'. More particularly it is used herein as defined in the Fibre-Channel or the Infiniband Standards.
- [0026] Thus, in terms of network topology, it can be considered as an 'agglomeration, i.e., a 'cloud'-like structure of point-to-point connections', in which the bandwidth

availability is less restricted compared to conventional bus technology, for example.

[0027] As a primary advantage the option is provided to offer a variable number of communication channels without the provision of a large local memory with a fixed size and performance in the network coupling means itself. Thus, the performance of said network coupling means can be easily scaled up according to dynamically changing traffic load without adding large amount of fast and expensive SRAM/DRAM area locally into said device. Thus, the traffic load flexibility is increased significantly.

[0028] The above inventional concept can be advantageously used with InfiniBand technology because a modern industry standard is defined therewith which allows to apply said fabric-based concept in the whole range of applications as it was sketched out above.

[0029] When the transmission control comprises the processing of address translation, e.g., logical to physical and vice versa- and protection information, e.g., tables then prior art remote processes can be advantageously performed.

[0030] When used for connecting a plurality of I/O devices associated with one or more computing devices as described above, the entire I/O periphery can be controlled with less restrictions and better performance according to the invention.

[0031] When said transmission control information is bundled per queue or queue pair then the number of cache line transfers to said interconnected memory means for a queue work request is reduced which increases performance and saves bandwidth.

[0032] Said cache memory can be configured for special queues not to discard transmission control information after cast out, i.e., after copying said control information back to the main memory. Subsequent repeated cast-in operations can be avoided if the cache line has not been re-used for other control information. A reduced number of cast-in operations and reduced latency improve the processing of queue pairs.

[0033] When writing said transmission control information to the local memory only before signaling the completion of a InfiniBand verb, then bandwidth is saved as well.

[0034] The inventional method can even be used for providing interprocess

communication (IPC) between a plurality of processes associated with one or more computing devices, independent of the underlying hardware structure of the network.

[0035] Furthermore it can be advantageously combined with the InfiniBand Architecture specification which was recently published. Amongst the general understanding of a person skilled in the art the following terms are thus used with additional – but not restricted to – particular reference to the InfiniBand Architecture specification, as it was recently published: "Adapters" in the sense of Host Channel Adapters (HCA) or Target Channel Adapters (TCA), "network" including a fabric, or "verbs" which provide an abstract definition of the functionality provided to a host by a Host Channel Interface (HCI).

[0036] The basic idea of the invention is, to use system memory as well as host adapter local memory for the transmission control information. The local memory is used like a cache, the system memory hold those work requests, which do not fit into the cache. This allows to provide the performance of implementing local memory only, but overcomes the resource restrictions of the local-memory-only approach.

[0037] The InfiniBand Architecture is designed around a point-to-point, switched I/O fabric, whereby end-node devices (which can range from very inexpensive I/O devices like single chip SCSI or Ethernets to very complex host computers) are interconnected by cascaded switch devices.

[0038] The invention provides a general means for improving prior art hardware-to-hardware data transmission on a very large range of scales: the invention can thus be advantageously applied to improve data traffic in pure, dedicated network devices like switches and routers, and furthermore, it can be well applied within LAN/WAN-based interprocess communication. The invention's basic concept is open to integrate any prior art network technology and in particular it can be advantageously applied to techniques such as Ethernet or Fibre Channel.

[0039] Thus, according to the invention any hardware-based data transmission like a module-to-module interconnection, as it is typified by computer systems that support I/O module add-in slots or chassis-to-chassis interconnections as they are typified by interconnecting computers, external storage systems or even external LAN/WAN

access devices, such as switches, hubs and routers in a data-center environment can be advantageously supported by the invention's concepts.

[0040] In Fig. 2 the system memory 12 – depicted left – has a plurality of entries 22 for storing the work request related control information for a particular work queue in a queue pair control block 22, further abbreviated as QPCB, each entry comprising a storage field 24 for storing the control information for it. Other queues are managed in here as well, these are, however not depicted in order to improve clarity.

[0041] Among others, the QPCB 22 comprises the following basic data:– queue pair state information– sequence numbers– maximum transfer unit size– destination LID (Local Identifier of connected queue pair)– destination GID (Global Identifier of connected queue pair)– error counters– performance countersAmong others the control field 24 comprises the following basic data:– send and receive queue head and tail pointers– number of associated completion queues– depth of send and receive queuesIn the host adapter memory 20 several transmission control blocks, e.g., a protection table PT, the work request queue WQ, with the queue pairs QP, an interrupt queue IQ, and a complete queue CQ, are managed. For each queue a plurality of cache entries 26 is provided for receiving the queue pair ID, i.e., a unique number and the respective control information required for the connecting host adapter to do its job, i.e., route the requested data to the correct network element or I/O device, respectively.

[0042] Further, a n-way associative array 32 is provided for storing the queue pair number 34 with the local address 36 in the cache storage 20, like it corresponds to usual caching techniques. Further, a QPCB directory 30 is provided for storing the queue pair number with the address of the system memory 12, in order to enable for casting out an entry from the cache memory 20 back into the system memory 12, when required.

[0043] In Fig. 2 enough local memory space is available in the QP area. During operation of the caching mechanism a situation emerges in which there is no free entry in said storage area for the queue pairs. This is depicted in fig. 3 which has basically the same structure as described above with reference to fig. 2.

[0044] With general reference to the figures and with special reference now to fig. 4 A

and B, the operation of the proposed caching technique will be described in more detail with a sample queue pair as it is defined in the InfiniBand Architecture: a send queue and a receive queue. It should be noted that any other queues required for compatibility with the InfiniBand Architecture, for example, or with other protocols can be managed according to the same principle.

[0045] On execution of a CreateQueue verb, e.g., when a queue pair shall be created, step 410, this is initiated by the CPU-memory-subsystem 8 in fig. 1. The respective application which originates the queue pair generation thus triggers that a queue pair control block (QPCB) is built in the system memory 12, step 420.

[0046] Then the host adapter's cache memory gets a request for storing caching data for the queue pair, i.e., the host adapter 18 gets a door bell signal indicating that the control area 24 of the new control block has to be copied, step 425 to the host adapter. A control logic decides, step 430, if enough free storage space is available in the cache memory. If not, see the NO-branch 430 of fig. 4A, then a classical cast-out / cast-in process takes place: In this situation now the host adapter checks the available storage space and detects that the local cache memory 20 is out of free space. Thus, in a next step 440 one particular queue pair, i.e., only its control information, is cast-out from the local cache memory according to an algorithm, like for example used in conventional caching techniques where for example the least recently used cache entry is overwritten (LRU algorithm). Thus, this entry is written back into system memory, step 450, and the address of the QPCB 22 is saved in the QPCB directory 30, step 460.

[0047] Then, in a next step 470 the host adapter 18 writes the new queue pair control block into the respective storage location, for example by simply overwriting the former contents of it.

[0048] Finally, the cache directory 32 is updated again, step 480. Then, the host adapter 18 is enabled to process the new queue pair, step 490.

[0049] When enough space is available in the local cache memory, see the YES-branch of decision 430, then the sequence of steps for cast-in/cast-out is not required. Instead, see fig. 4B now, the respective control information is copied from the system memory

12 to the local memory 20 of the host adapter 18, step 510. Thus, only a small fraction of the queue pair data amount, i.e. only the control information is stored in the local cache memory 20.

[0050] Further, said cache directory 32 is updated, as it would be done with usual caching techniques known in prior art within a processor unit, step 520. Then the request is ready for execution, the queue pairs can be processed, step 530.

[0051] Thus, the invention represents a large step forward to a significantly increased performance in host adapter's work request handling because all transmission control information – which requires only small chip area compared to the total work request data contained in the queue pair – is available immediately where it is required: local to the host adapter. The rest of data which can be sent "through" the host adapter is stored external to the adapter/switching element because it does not carry any routing/switching information. Thus, a person skilled in the art will appreciate that the inventional concept can be scaled up and down easily with a small increase or decrease of required chip area needed for the local cache memory 20 – according to the actual requirements present on a given hardware and traffic situation.

[0052] In the foregoing specification the invention has been described with reference to a specific exemplary embodiment thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are accordingly to be regarded as illustrative rather than in a restrictive sense.

[0053] For example, the way in which the cache memory 20 is operated, can be varied to the different types known in the art, e.g., write back, or write through, etc..